

(19)



(11)

EP 3 961 401 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
08.11.2023 Bulletin 2023/45

(51) International Patent Classification (IPC):
G06F 11/16 ^(2006.01) **G06F 11/20** ^(2006.01)
G06F 11/14 ^(2006.01)

(21) Application number: **21193085.4**

(52) Cooperative Patent Classification (CPC):
G06F 11/1633; G06F 11/1407; G06F 11/1675;
G06F 11/202; G06F 11/2033; G06F 11/2064;
G06F 2201/815; G06F 2201/82

(22) Date of filing: **25.08.2021**

(54) REAL-TIME FAULT-TOLERANT CHECKPOINTING

FEHLERTOLERANTES CHECKPOINTING IN ECHTZEIT

POINTAGE DE CONTRÔLE TOLÉRANT AUX FAUTES EN TEMPS RÉEL

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR

(30) Priority: **26.08.2020 US 202017003808**

(43) Date of publication of application:
02.03.2022 Bulletin 2022/09

(73) Proprietor: **Stratus Technologies Ireland Limited**
Dublin 1 (IE)

(72) Inventors:
• **HORVATH, Charles J.**
Jamaica Plain, 02130 (US)
• **CAO, Lei**
Westford, 01886 (US)

(74) Representative: **Purdylucey Intellectual Property**
6-7 Harcourt Terrace
D02 FH73 Dublin 2 (IE)

(56) References cited:
WO-A1-2012/084839 WO-A1-2016/054770
US-A1- 2007 174 484

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 3 961 401 B1

Description

BACKGROUND

[0001] Checkpointing technology is a means to achieve fault tolerant behavior of computer systems. Current checkpointing technology, either virtualization-based or OS-based, is incompatible with hard real-time operation in that a checkpoint can be required to occur at any point in time, thereby impacting the system response time and causing a missed a deadline.

[0002] Applications and systems requiring both real-time support and fault-tolerant operation are currently designed with custom software and/or hardware. They are bespoke and tailored for particular industries and purposes. This requires significant expense and time to develop, deploy, and maintain such a system. US Patent Publication No. US 2007/174484 describes an apparatus and method for high performance checkpointing and roll-back of network operations. International Patent Publication No. WO 2016/054770 describes techniques for checkpointing/delivery between primary and secondary virtual machines. International Patent Publication No. WO 2012/084839 describes a method and system for virtual machine failover management.

[0003] The disclosure provides systems and methods of providing real-time fault tolerance suitable for use in various industries and support levels, while addressing the challenges associated with checkpointing noted above.

SUMMARY

[0004] In part, the disclosure relates to a real-time fault tolerant computing system and a method of controlling an operational system in real-time as set out in the appended claims. The system includes a first computing device, a second computing, and a hardware interconnect. The first computing device includes one or more memory devices, one or more processors, a first network interface operable to receive device data and transmit output data over a time-slot-based bus, wherein the output data is generated from processing device data, and a first real-time checkpoint engine. The second computing device includes one or more memory devices, one or more processors, a second network interface operable to receive device data over the time slot-based bus, and a second real-time checkpoint engine. In one embodiment, the hardware interconnect is operable to permit data exchange between the first computing device and the second computing device. Checkpoints are generated by the first or the second real-time checkpoint engines during lower-priority communication time slots allocated on the time slot-based bus to avoid interfering with any real-time communications to or from the first and second computing devices.

[0005] In one embodiment, if the first computing device experiences a failure the second computing device con-

tinues processing device data to generate the output data and transmit the output data using the time slot-based bus. According to the invention, checkpoint operations are always postponed until the occurrence of a lower-priority communication slot, even upon the first computing device and the second computing device producing divergent output data prior to the lower-priority communication slot occurring. According to the invention, the first computing device and the second computing device both delay release of output data to one or more devices in communication with the time-slot based bus until after a subsequent checkpoint to avoid release of divergent output data should a failure occur.

[0006] In one embodiment, the first computing device generates checkpoints on a per cycle basis when output data includes low priority data is transmitted using a low priority time slot. In one embodiment, the device is a sensor and the output data includes sensor data, the sensor data is operable as an input for a control system. In one embodiment, the first computing device and the second computing device are both actively processing device data and producing output data so that if the first computing device experiences one or more failure modes a time delay for second processing device to continue processing is about zero or zero.

[0007] In one embodiment, the duration of one or more time slots is increased by a detection period, wherein the detection period ranges from about 1 microsecond to about 100 microseconds. In one embodiment, the first computing device and the second computing device both delay release of output data to devices in communication with the time-slot based bus until after the next checkpoint to avoid release of inconsistent output data. In one embodiment, if the first computing device is operating normally after a checkpoint, the first computing device releases any delayed output data. In one embodiment, the first computing device transmits all of its output data to the second computing device to avoid such output data being dropped upon occurrence of a failure of the first computing device.

[0008] In one embodiment, if the first computing device fails, the second computing device takes over release of the first computing device's output data that was generated during the last checkpoint. In one embodiment, only the first computing device handles release of output data when no failures have occurred and in parallel shares such output data with the second computing device. In one embodiment, the device data from one or more devices are transmitted to the first computing device and the second computing device over the time slot-based bus. In one embodiment, the first computing device and the second computing device receive incoming data directly through a broadcast or a multicast.

[0009] In one embodiment, each of the first computing device and the second computing device include a respective non-virtual operating system (OS) that includes a respective network driver, checkpoint shim, and storage driver stored in the one or more memory devices and

executing on each of the one or more processors. In one embodiment, the network driver includes a replicator in communication with a release coordinator and a release notifier. In one embodiment, the one or more networking programs stored in the one or more memory devices of the first and second computing devices delay and coordinate transmission of output data using one or more of the hardware interconnect, the first network interface and the second network interface. In one embodiment, each of the first computing device and the second computing device include a respective virtual operating system (OS) that includes one or more respective network support applications, a checkpoint engine, and one or more storage support applications stored in the one or more memory devices and executing with respective hypervisors. In one embodiment, checkpoints are transmitted using the hardware interconnect.

[0010] In part, the disclosure relates to a method of controlling an operational system in real-time. The method includes receiving device data, at a primary node and a secondary node, from one or more devices in electrical communication with a time slot-based bus, wherein the bus is operable to receive data according to time slots of two or more priority levels; during a first time slot, generating a checkpoint at primary node and transmitting the checkpoint data to secondary node, wherein the first time slot is a lower priority time slot; generating an output data at the primary node and the secondary node in response to processing the received device data; and transmitting, from the primary node, the output data over the bus using a second time slot, wherein the second time slot is a higher priority time slot, wherein both the first node and the second node are in an active mode.

[0011] In one embodiment, the method further includes transmitting data during a non-checkpoint portion of a lower priority time slot. In one embodiment, the method further includes gating network data using a respective network software module executing on both the primary and the secondary node such that output data is released from the node whose state has carried forward to current checkpoint. In one embodiment, the method further includes synchronizing transmission of output data with high priority time slots, wherein the output data is received by the operational system.

[0012] In one embodiment, the method further includes transmitting data during a non-checkpoint portion of a lower priority time slot. In one embodiment, the method further includes using the primary node, all of its respective outgoing data to the secondary node to avoid such outgoing data from being dropped upon occurrence of a failure of the primary node. Checkpoint data is transmitted using a hardware interconnect.

[0013] Selectively releasing data such that it is synchronized with a time slot having a priority level that matches that or is otherwise compatible with that of the data to be transmitted using a timeslot based-bus is performed in various embodiments of the disclosure. In addition, the transmission of data using the hardware inter-

connect that is processor / computing device intensive such as, for example, one or more stages of a checkpoint are transmitted during low priority timeslots on the timeslot-based bus. In various embodiments, the hardware interconnect is a separate communication channel from the timeslot-based bus.

[0014] Although, the disclosure relates to different aspects and embodiments, it is understood that the different aspects and embodiments disclosed herein can be integrated, combined, or used together as a combination system, or in part, as separate components, devices, and systems, as appropriate. Thus, each embodiment disclosed herein can be incorporated in each of the aspects to varying degrees as appropriate for a given implementation.

BRIEF DESCRIPTION OF DRAWINGS

[0015] The structure and function of the disclosure can be best understood from the description herein in conjunction with the accompanying figures. The figures are not necessarily to scale, emphasis instead generally being placed upon illustrative principles. The figures are to be considered illustrative in all aspects and are not intended to limit the invention, the scope of which is defined only by the claims.

FIG. 1A is a schematic diagram of an exemplary cyclic real-time system that may be implemented using the fault tolerant features disclosed herein.

FIG. 1B is a schematic diagram of various devices in communication with a real-time fault tolerant system that shares a bus with various devices in accordance with an embodiment of the disclosure.

FIG. 2 is a schematic diagram of real-time fault tolerant system showing the path of real-time device data and output data in accordance with an embodiment of the disclosure.

FIG. 3 is a schematic diagram of real-time fault tolerant system that includes a first computing device and a second computing device in accordance with an embodiment of the disclosure.

FIG. 4 is a schematic diagram showing checkpoint coordination relative to various timeslots of a shared bus for use in a real-time fault tolerant system in accordance with an embodiment of the disclosure.

FIG. 5 is a schematic diagram showing three cycles of operation for a primary and a secondary RT/FT real-time fault tolerant implementation in which packets corresponding to output data are held or gated. FIG. 6 is a schematic diagram showing three cycles of operation for a primary and a secondary RT/FT real-time fault tolerant implementation that shows real-time checkpointing in accordance with an embodiment of the disclosure.

FIG. 7 is a schematic diagram showing three cycles of operation for a primary and a secondary RT/FT real-time fault tolerant implementation.

FIG. 8 is a block diagram of a real-time fault tolerant virtualization-based checkpointing system according to an embodiment of the disclosure.

FIG. 9 is a block diagram of a real-time fault tolerant operating system-based checkpointing system according to an embodiment of the disclosure.

FIG. 10 is a block diagram of various networking components suitable for use in a virtualization-based real-time fault tolerant system in accordance with an embodiment of the disclosure.

FIG. 11 is a block diagram of various networking components suitable for use in an OS-based real-time fault tolerant system in accordance with an embodiment of the disclosure.

FIG. 12A and 12B are schematic diagrams showing four cycles of operation for a primary and a secondary RT/FT real-time fault tolerant computing device implementation in accordance with an embodiment of the disclosure.

DETAILED DESCRIPTION

[0016] In part the disclosure relates to various systems and methods of real-time fault tolerant computing. A real-time computer system provides guaranteed response times, that is time from an input or inputs to an output or response from the computer system is guaranteed by system design not to exceed a predetermined amount. This predetermined amount of time is often called a deadline. There are multiple levels of real-time support, including hard real-time support and soft real-time support. Various embodiments of the disclosure implement real-time checkpointing in concert with a timeslot-based bus to process device data, such as from sensors and other devices such as computing devices, and generate output data to control actuators or otherwise generate inputs for control systems and other systems, devices, and components as may be applicable for a given use case or implementation.

[0017] In the context of hard real-time support, missing a deadline results in a total system failure. Hard real-time support is used when a failure to meet a specified deadline can lead to catastrophic consequences. Examples include, without limitation, chemical plant processes, nuclear power station control, railway signal control, avionics fly-by-wire, and automotive engine control. Conversely, in the context of soft real-time support, the usefulness of a result degrades after its deadline, thereby degrading the system's quality of service. Examples include, without limitation, a computer sound system for which missed deadlines may cause some audio distortion, a computer gaming system for which missed deadlines may cause some delays in player action, or telephone switch systems for which a missed deadline may result in failure to initiate a call. The various systems and methods disclosed herein may be used for either hard real-time support or soft real-time support.

[0018] This disclosure describes a checkpointing

based solution for achieving application transparent real-time fault-tolerant operation which can run on standard (commodity) computer systems. Use of this system achieves significant cost and time savings as compared to purpose designed solutions with custom software and/or hardware. In part, the disclosure relates to various checkpointing based systems which provides application transparent hard real-time fault-tolerant operation which can run on standard (commodity) computer systems. These systems offer various enhancements relating to checkpointing.

[0019] Conventional checkpointing is virtualization-based. Specifically, it is a means to achieve virtual machine (VM) level fault tolerant operation by periodically stopping a VM, capturing its state, and replicating the captured state on a second VM so that the second VM can continue operation should the first VM fail. The process of stopping a VM, capturing its state, and replicating the state to another VM is called a checkpoint. Checkpointing VMs is called virtualization-based or hypervisor-based checkpointing.

[0020] It is also possible to implement checkpointing at the OS level, i.e. in OS drivers rather than in a hypervisor. The various, systems, methods, combinations thereof, and other embodiments described in this disclosure applies equally well to virtualization-based and OS-based checkpointing.

[0021] Generally, improving upon checkpointing-based approaches in fault tolerant systems along with the use of timeslot prioritization for device data and selecting low priority timeslots for checkpointing are some of the features used to implement real-time fault tolerant systems and methods. With the foregoing as a high-level introduction, it is useful to consider checkpointing in more detail.

[0022] Checkpointing is a technique to keep a pair of virtual machines (VMs) or operating systems (OSs) at sufficiently the same state so that if one VM or OS fails the other VM or OS can continue the failed VM's or OS's operation essentially transparently to client systems. To achieve non-stop service, also called fault tolerance (FT), checkpointing deploys the pair of redundant VMs or OSs on separate physical servers so that if one physical server fails, the VMs or OSs and applications continue to run on the remaining physical server.

[0023] Checkpointing is necessitated by the fact that modern multi-threaded or multi-core computer systems do not execute deterministically. Two computers receiving the same input data will not execute the same instruction streams and hence may produce different output data (or output data timing) due to various factors. For example, hardware interrupts not being received at the exact same point in the instruction stream on both computers may ultimately cause processes to be scheduled differently between the two computers. In addition, disk response time differences (rotational latency variations for rotational drives, and garbage collection/remapping activity variations for solid-state drives) may ultimately lead

to processes being scheduled differently on the two computers. Furthermore, general non-deterministic processor behavior such as thermal throttling, different cache hit/miss behavior due to earlier instruction/data accesses, non-precise alignment of fine granularity clocks and counters, etc., may ultimately lead to processes being scheduled differently on the two computers. Checkpointing periodically brings a pair of (non-deterministic) OSs or VMs back to the same effective state, so should an OS or VM fail, the partner is able to resume operation for the other transparently to external client systems.

[0024] At a high level, checkpointing periodically stops a running VM or OS, captures its state, and sends the state (or state changes) to a partner VM or OS, keeping the partner VM or OS at a near equivalent state. In the case of virtualization-based checkpointing the checkpointing functionality is largely implemented within a hypervisor. In the case of OS-based checkpointing the checkpointing functionality is largely implemented within operating system drivers. The time a checkpointed VM or OS is running is called the *run interval* and the time a checkpointed VM or OS is stopped to capture state is called the *checkpoint*. The sum of these two intervals is called the *epoch*. Exemplary checkpoint epochs and run intervals are showed in the context of scheduling data for release on a TSB in FIGS. 4, 12A and 12B, which are discussed in more detail herein.

[0025] In the case of OS-based checkpointing the OS is largely stopped during the checkpoint, but portions of the OS associated with the checkpointing process continue to run during the checkpoint. Various embodiments of the disclosure may improve upon and/or leverage some of the features disclosed in co-pending application no. 16/900,909, entitled "OPERATING SYSTEM-BASED SYSTEMS AND METHOD OF ACHIEVING FAULT TOLERANCE", filed on June 13, 2020.

Real-Time Overview

[0026] Real-time computer systems provide guaranteed response times, that is time from an input or inputs to an output or response from the computer system is guaranteed by system design not to exceed a predetermined amount. This predetermined amount of time is called a deadline. FIG. 1A is a schematic diagram of an exemplary cyclic real-time system 5 that may be implemented using the fault tolerant features disclosed herein. Many real-time systems operate in a cyclic fashion, that is inputs are received and outputs are supplied at a repeating fixed interval. That fixed interval is called the cycle time (CT). The real-time fault tolerant system 7a receives sensor inputs such as from various devices, including sensors, and generates outputs such as control outputs. The control outputs can be used to control various devices such as actuators and other controllers. The real-time fault tolerant system 7a is operable to implement real-time checkpointing using a virtualized approach or an operating system (OS) based approach.

[0027] In industrial and process control environments, the communication medium connecting a controlling computer system and the sensors, actuators, valves, and other entities is often a shared bus. As shown in FIG. 1B, the overall manufacturing system 12 has various devices that are in communication with a shared bus. The shared bus shown is implemented using a timeslot-based bus (TSB). The various devices D1 through D7 correspond to a temperature sensor (D1), a gauge (D2), a valve (D3), a motor (D4), a robotic arm (D5), a conveyor belt (D6), and an alarm (D7). Devices D1 and D2, the temperature sensor and gauge, may serve as sensors that provide sensor inputs to the RT/FT system 7b. As shown, the system 7b typically includes a first computing device 8a and a second computing device 8b that are connected to each other and the TSB. The valve D3, motor D4, robotic arm D5, conveyor belt D6 and alarm D7 may be controlled based on the control outputs (output data) generated by the system 7b in response to the sensor inputs from devices D1 and D2. In turn, the alarm D7 may also serve as an input and an output with the alarm state of device D7 being fed back to the system 7b to regulate other device or to continuously indicate that everything is working according to acceptable levels.

[0028] Many real-time control systems use a shared communication medium such as Ethernet for the TSB embodiment shown in FIG. 1B. For simplicity of presentation the TSB in FIG. 1B is illustrated as connecting directly to devices (D1 to D7), in practice there is often an intervening electronic component between the device and field bus, typically a PLC (programmable logic controller). Various edge computing devices may also be connected to TSB in various embodiments. The real-time fault tolerant system 7b is operable to implement real-time checkpointing using a virtualized approach or an operating system (OS) based approach. The computing devices, such as a first computing device 8a and a second computing device 8b include network connections such that both devices receive device data through the TSB from one or more (or all) of the devices D1-D7.

[0029] To maintain real-time behavior with a shared bus, processes are implemented to prevent various computing devices that access the bus from simultaneously communicating on the shared bus. In part, the disclosure relates to protocols and other safeguards to prevent simultaneous access to the shared bus or otherwise sending data such as checkpoint data during high priority timeslots. Failure to implement such safeguards results in non-deterministic communication delays which in turn could cause real-time deadlines to be violated.

[0030] In various embodiments, avoiding communication conflicts on a shared medium such as a shared bus is to time-slice the communication medium. In various embodiments, the real-time fault tolerant systems disclosed herein allocate pre-assigned timeslots within a processing loop cycle or other period or cycle for each communicating entity on the shared bus. Examples of time-sliced communication buses include, without limita-

tion, ProfiNET IRT, CC-Link IE Field, Foundation Fieldbus, and Ethernet POWERLINK. Various TSBs are operable to send and receive data according to time slots having varying priority levels such as a high, medium, and low priority. Other more general priority levels may be used to manage how device data is scheduled relative to checkpoint data. A first priority level, a second priority level, a third priority level, etc. may be used to distinguish time slots to facilitate real-time checkpointing and avoiding conflicts that may result in errors that prevent real-time operation.

[0031] In various implementations, a timeslot-based bus is used to implement various real-time fault tolerant systems that collect data from sensors or other devices that have the capability to allocate specific timeslots for time critical communication and other time slots for non-time critical or less time critical communication. This permits the use of a shared communication medium fieldbus for non-deterministic communication without interfering with time critical communication on the same shared bus such as a TSB.

[0032] For some TSBs, a low priority time-sliced communication bus slots may be characterized as acyclic, meaning the communication during those slots is not constrained by cycle time driven deadlines. Various shared buses suitable for implementing the embodiments of the disclosure may include combinations of cyclic and acyclic features and timeslots such as timeslots having various levels of priority.

Time Sensitive Networking (TSN) Scheduled Traffic

[0033] As an exemplary TSB is informative to consider TSN (Time Sensitive Network) technology that may be used to implement various TSB configurations. TSN supports several methodologies for achieving bounded low-latency network data transport. These are now incorporated into IEEE 802.1Q-2018. Examples of TSN features may include one or more of the following Credit Based Shaper (formerly 802.1Qav), Preemption (formerly 802.3br & 802.1Qbu), Scheduled Traffic (formerly 802.1Qbv), Cyclic Q-ing & Fwd (formerly 802.1Qch), and Async Shaping (formerly P802.1Qcr).

[0034] Various embodiments of the disclosure may use suitable TSBs such as TSN scheduled traffic (IEEE802.1Qbv, a.k.a. TSN time-aware shaper, a.k.a. TSN time-aware scheduler) to illustrate real-time behavior. TSN is an example of a TSB suitable for use in various embodiments.

[0035] The TSN time-aware scheduler is a time-sliced cyclic bus based on ethernet functions. TSN is an ethernet based non-proprietary standard. References to using TSN for a shared-bus for the real-time systems and methods disclosed herein are provided as an illustrative example. The disclosure is not specific to TSN based networks and may work with any time-sliced cyclic bus as well as with other non-cyclic real-time solutions for which there are periodic time slots where a checkpoint can be

taken without interfering with real-time activities. In general, embodiments of the disclosure may be implemented using a time-slot-based bus (TSB) such as for example a TSN-based bus or network and others disclosed herein or suitable for providing real-time functionality as disclosed herein.

Checkpointing and Real-Time Operation

[0036] In various aspects, the disclosure relates to real-time checkpointing and its role in supporting real-time fault tolerance. Existing legacy checkpointing solutions may be grouped into two classifications, active-passive checkpointing and active-active checkpointing. These existing legacy checkpointing technologies are described in more detail below. Real-time checkpointing is a different type of checkpointing than these legacy technologies. Real-time checkpointing addresses the shortcomings of active-passive checkpointing or active-active checkpointing which makes them unsuitable for real-time operation. Real-time checkpointing is a different type of checkpointing than active-passive checkpointing or active-active checkpointing. Real-time checkpointing addresses the shortcomings of active-passive checkpointing or active-active checkpointing and hence is suitable to real-time operation.

[0037] A high level illustrative example of a real-time system 7c is shown with a TSB as part of an overall control system 20. In that example, shown in FIG. 2, sensors 1 and 2, S1, S2, are providing real-time (RT) data to the real-time fault tolerant (RT/FT) checkpointing based system, labeled RT/FT, and the RT/FT system is processing that data to control the actuator A in real-time.

[0038] The sensors S1, S2 in the diagram can be a variety of devices (pressure sensors, flow sensors, speed sensors, fluid level sensors, proximity sensors, etc.) and have a variety of implementations. For example, the various devices D1 through D6 in system 5 of FIG. 1B may include various sensors and other devices. Various sensors and devices, suitable for use with the systems and methods disclosed herein may include, without limitation, a smart networked sensor; a networked PLC (programmable logic controller) connected to a sensor or sensors; a computer hosting virtualized PLCs connected to sensors; distribute Control Nodes (DCNs, industrial control devices with IO and compute capabilities); combinations thereof, and others as disclosed herein or otherwise applicable to a given company's products and services. Similarly, the actuator can be a variety of device types and have a number of possible implementations.

[0039] If the sensors were position, proximity, and/or pressure devices associated with a robotic arm and the actuator represented the robotic arm's servo motors controlling the its movement, the sensors would be supplying data in real-time to the RT/FT system which is processing by the RT/FT to control the robotic arm's motions via the arm's servo motors.

[0040] Still referring to FIG. 2, the two exemplary de-

vices S1, S2 are shown along with an actuator A. In various embodiments, devices S1 and S2 are sensors that collect data relative to a system or process that is being monitored or controlled by the control system 20. Real-time device data from the sensors is being processed by the RT/FT 7c during one or more processing cycles and then relayed to the actuator as real-time control information. A given RT/FT may include a network interface to receive device data (such as sensor data) and transmit output data. The output data is generated from processing received device data and optionally other inputs and parameters that are specific to the process or system that is being controlled or monitored. The device data and output data are typically transmitted using a TSB. In the system of FIG. 2, the RT/FT is processing the data from S1 and S2 to generate output data such as real-time control information that will be transmitted to the actuator A to cause it to respond in a particular manner or perform a given task or take a particular action in response to the received output data. The RT/FT 7c is simplified in terms of the processing details and how one or more computing devices are configured to provide fault tolerance and a real-time implementation. FIG. 3 shows an exemplary RT/FT with more details and includes two computing devices.

[0041] The RT/FT system 7d of FIG. 3 includes two computing devices that may be identified as primary and secondary (or vice versa) or RT/FT¹ and RT/FT² (or vice versa). The RT/FT system 7d is part of an overall control system 30 that includes a TSB, S1, S2, and A. The devices / sensors S1 and S2 and the actuator A are the same as depicted in the system of FIG. 2, but additional details of their connections to the primary and secondary computing devices are shown. The two computing devices work in concert to implement the fault tolerant real-time checkpointing based system 7d. These computing devices may also be identified as first and second computing devices, servers, nodes, or other terms as disclosed herein. Notwithstanding the foregoing, a given first computing device may initially have a first role while the second computing device may have a second role. The first and second roles may switch or be changed to other roles if one of the computing devices experience an error or a failure mode.

[0042] The system of FIG. 3 includes a shared bus, such as a TSB, and also may include network connections to a network such as an information technology (IT) network of an enterprise, the internet, or other networks. Each of the computing devices RT/FT¹ and RT/FT² are connected by a hardware interconnect HIC. The sensors, actuator(s) and RT/FT systems are connected using redundant shared networks 17 represented by the thin dark lines. An exemplary data flow is illustrated by the wide light lines 18. The various arrow heads show the direction of data flows.

[0043] For the non-failure case corresponding to normal operation without a failure in 7d or one of the two computing devices, the dataflow is as follows. The sensor

/ device data is received and processed by both RT/FT nodes (so that if one node fails no processing time or state is lost). Thus, if RT/FT¹ were to fail, RT/FT² would be in receipt of the same sensor / device data and be able to generate the output data in order for the control system 30 and the components that need real-time outputs to maintain real-time operation such as actuator A. In one embodiment, the primary node (RT/FT² in the above diagram) supplies the control output / output data to the actuator(s) A. This is an exemplary requirement. In other instances, the secondary node (RT/FT¹ in the above diagram) supplies the control output / output data to the actuator(s) A. In another embodiment, both the primary computing device and the secondary computing device can generate the output data and transmit the output data redundantly with the actuator operable to be able to process the duplicative signals and ignore duplicative inputs in some embodiments.

20 Checkpointing

[0044] In various embodiments, checkpoints are coordinated with a TSB so that checkpoints overlap or a synchronized for release only with the lower-priority (non-real-time) communication timeslots of the TSB. Coordinating and scheduling checkpoints and checkpoint data using lower priority timeslots prevents a checkpoint operation from interfering with real-time communication. For example, if checkpoint data was competing with output data needed by an actuator to maintain a rate of fluid flow in a chemical engineering plant, it could be disastrous if checkpoint data was sent using a timeslot needed to control the actuator. FIG. 4 is a schematic diagram showing various scheduling interactions relative to timeslots of an exemplary TSB and checkpointing operations. In general, FIG. 4 shows an exemplary approach to coordinating checkpoints of a RT/FT system with a TSB.

[0045] Specifically, as shown in FIG. 4, a RT/FT system is processing device data from sensors and generating output data for actuators on a per cycle basis. Two processing cycles of the RT/FT are shown TSB Cycle 1 and TSB Cycle 2. A given TSB cycle includes N time slots. In the example shown for each cycle, four time slots are shown, thus N=4. The duration of the time slots may be the same or they may be different. As shown, time slot 4 is allocated for low priority traffic and has a longer duration than time slots 1 - 3. Each of the respective cycles spans two parallel data generating events. These events include checkpoint scheduling and TSB network scheduling. As noted above, the time a checkpointed VM or OS is running is called the run interval and the time a checkpointed VM or OS is stopped to capture state is called the checkpoint. The sum of these two intervals is called the epoch.

[0046] A first checkpoint epoch and a second checkpoint epoch with their own respective run interval and checkpoint during which the VM or OS is stopped during time slot 4 are shown in FIG. 4. The first epoch corre-

sponds to the TSB Cycle 1 and the second epoch correspond to TSB Cycle 2. These respective time periods may differ in some embodiments. As shown, after the run interval for a given computing device of a RT/FT, a checkpoint is released such that it overlaps with or is synchronized with a time slot, such as time slot 4 in the example shown for each epoch. Fixed-cycle checkpoints are scheduled to occur during TSB (or TSN as shown in this example) during low-priority time slot on the TSB to avoid interfering with real time traffic. Various examples of real-time traffic are shown with regard to time slots 1, 2 and 3 for each TSB cycle.

[0047] Another element of the disclosure is that both of the computing devices of the RT/FT are active. In contrast, in some fault tolerant system, one computing device is passive and needs time to effectively take over from the active computing device to which it is paired. By having both nodes active, no processing time or state is lost when/if the primary fails and hence real-time deadlines can be maintained even when the primary node fails. In addition, to support real-time fault tolerance network packets are gated until after the next checkpoint. FIG. 5 is a schematic diagram showing three cycles of operation for a primary and a secondary RT/FT real-time fault tolerant implementation in which packets corresponding to output data are held or gated. FIG. 5 offers more detail relative to the embodiment of FIG. 4 with each of the computing devices (primary RT/FT and secondary RT/FT) requiring time slot management over multiple processing cycles (cycle 1, cycle 2, and cycle 3 are shown but any number of cycles is possible). In turn, during each cycle, a run interval and a checkpoint is also occurring. Thus, in various embodiments, for N cycles, there are N epochs, absent a failure or some other event.

[0048] As shown in FIG. 5, network packets are gated or held because during the run interval the two nodes can diverge. During each of the run intervals, outgoing network packets that would be sent over the TSB are held. If a failure were to occur as part of the RT/FT and packets were being released immediately, the packets earlier released by the now failed node may not be state consistent with the packets released by the former secondary now primary node after the failure. By gating the packets, the RT/FT ensures the packets released were all generated by the node whose state is carrying forward from the previous epoch. In various embodiments, gating packets may also be referred to as delaying packets. Gated packets are released during the correct TSN time slot. If this not adhered to, a given real time system may cease to be fault tolerant and experience errors which in turn may prevent real-time operation. Outgoing packets from the previous run interval are released after a checkpoint has occurred during a low priority time slot.

[0049] The gated outgoing network packets from primary RT/FT are sent to secondary RT/FT. The secondary RT/FT generally operates to hold outgoing packets received from primary and discards those packets if the primary does not fail during the cycle the primary is re-

leasing its copy of the corresponding packets. In a non-failure cycle, the secondary RT/FT discards its state from that cycle and continues from the last primary checkpoint. In addition, in a non-failure case, secondary RT/FT's packets are discarded upon completion of checkpoint. The above points discussed with regard to FIG. 5 may be more clearly illustrated by examining a failure case.

[0050] FIG. 6 is a schematic diagram showing three cycles of operation for a primary and a secondary RT/FT real-time fault tolerant implementation in the case of a failure. As was the case with FIG. 5, there are three cycles divided into timeslots for scheduling device data and output data relative to the timeslots available on the TSB. In the first cycle, cycle 1, there is no failure and each of the primary RT/FT and the secondary RT/FT schedule device data and output data (for actuator) in time slots 1, 2 and 3, with low priority traffic corresponding to checkpoint handled on time slot 4.

[0051] As shown, prior to cycle 1, gated network packets from the previous cycle are sent from the primary to the secondary. During the run interval aligned with part of cycle 1, the secondary is running independently of the primary following the last checkpoint. During cycle 1, gated packets from the primary are sent to the secondary. In a non-failure case, the secondary discards its state and continues from last primary checkpoint.

[0052] As shown in FIG. 6, in cycle 2 there is a failure event at the primary RT/FT and cycle 3 shows details post-failure. When the primary fails during cycle 2, the secondary has gated network packets from the last checkpoint to release. In addition, the secondary is running based on a state that is consistent with released packets. As a result, the secondary's execution state is valid and no execution time is lost. Effectively, the secondary can pick up immediately where the primary left off with no loss of processing time. No checkpoints are needed after failure, but in some preferred embodiments, it is desirable to gate packets as if checkpointing occurred. In some non-preferred embodiments where packets are not gated as if checkpointing is still occurring, post failure the secondary may try to transmit twice as many packets during a critical (high priority) time slot from the last checkpoint and those from the current checkpoint. In other non-preferred embodiments where packets are not gated as if checkpointing is still occurring, post-failure the secondary may miss sending packets during a cycle when the primary returns to operation, which may break the real time requirement of the overall RT/FT. In cycle 3, after the failure of primary, the secondary continues running from its state and in preferred embodiment releases its network packets which are consistent with its processing state the following cycle.

Networking Features and Management of Outgoing Packets

[0053] Dropping outgoing network packets is not acceptable for true real-time operation. This follows be-

cause the time penalty for a higher-level protocol recovery would break real-time deadline commitments. Further, many real time busses assume reliable transmission for critical packets and do not have a higher-level recovery protocol. For example, TSN is at the ethernet level and with TSN reliable ethernet transmission can be achieved for critical packets since all collisions can be avoided.

[0054] FIG. 7 is a schematic diagram showing three cycles of operation for a primary and a secondary RT/FT real-time fault tolerant implementation. As illustrated in FIG. 7 below, the real-time checkpointing solution could drop outgoing packets on failover were not the mechanisms described below incorporated into the invention. Specifically, packets from the previous cycle (labeled cycle 1) have been queued by the primary RT/RT for release in the following cycle, labeled cycle 2. If the primary RT/FT fails during cycle 2, not all packets maybe be released. As shown in FIG. 7, during failure of primary in cycle 3, packets 40 are not transmitted because active node failure before it released all of the packets from the last epoch.

[0055] In part, the disclosure also relates to removing the possibility of dropped outgoing packets on failover. This may be achieved by the primary RT/FT sending its outgoing packets to the secondary, either during the run interval or at the checkpoint. Then, depending on the system constraints, one of two algorithms or methods can be applied for ensuring no outgoing packets are dropped when failure occurs. The first algorithm operates such that if the system can tolerate large numbers of duplicate packets, both the primary and secondary RT/FT can unconditionally release the outgoing queued packets during their assigned time slots. The second algorithm provides that if the system cannot tolerate large numbers of duplicate packets, the primary can supply the secondary and ongoing indication of which packets have been released. If the primary fails, the secondary would then resume releasing packets from where the primary left off prior to failure. Generally, packets are released during their assigned time slot.

[0056] It should also be noted that for the second algorithm above, allowance is made for the fact that there can be delay in the primary's release of a packet relative to providing an indication to the secondary that the packet has been released. One solution is to delay the indication of a packet's release to the secondary a sufficient amount to ensure the packet has in fact been released by the time the notification is received by the secondary. This does create a timing window wherein duplicate packets are a possibility, specifically at failover, a packet or some packets released by the primary can be re-released by the secondary.

Incoming Packets

[0057] For real-time checkpointing, both nodes are operable to receive incoming packets since both nodes are

actively running the OS and application workload. This is in contrast with active-active solutions which may operate with a primary VM receiving all incoming packets and forwarding them to the secondary so that they are received by both nodes. Such an active-active approach can result in dropped incoming packets. For example, the packets intended to be received between the time the primary fails and the network packets are being re-routed to the secondary, upon it becoming primary, will be dropped. Dropped packets are unacceptable for true real-time operation and thus an active-active implementation is not possible for a real-time implementation. Further, in various embodiments to avoid dropped incoming packets on failover, external transmitters of real-time network packets may use multicasts or broadcasts so that packets are being directly received by both nodes.

Overhead Considerations - Checkpoint Interval and Network Gating

[0058] Checkpointing introduces overhead associated with periodically stopping the OS or guest VM. There are limits to how quickly this can be done. Further, input to output latency is impacted by the need to gate outgoing packets across checkpoints. Depending on how the network time slots are assigned relative to the checkpoints, this added latency may be as short as the worst-case checkpoint interval or as long as a whole checkpoint cycle.

Fault Detection

[0059] In some embodiments, time-slot duration or period may be increased to account for fault detection time. For example, if during a timeslot, about 100 μ sec worth of data can be transmitted, the time allotted to that timeslot is increased by a detection time DT. The detection time DT allows for the time it takes a secondary to detect that a primary computing device has failed and resume transmission of the outgoing data. Absent the inclusion of a detection time DT, during a failure case not all outgoing data will be able to be transmitted during the timeslot. In various embodiments, the detection time DT ranges from about 0.1 microsecond to about 100 microseconds. In various embodiments, the detection time DT ranges from about 1 microsecond to about 500 microseconds.

Virtualization-Based Checkpointing

[0060] FIG. 8 is a block diagram of a real-time fault tolerant virtualization-based checkpointing system according to an embodiment of the disclosure. As shown in FIG. 8, a virtualization-based real-time FT checkpointing solution is implemented in fault tolerant system 70. As shown, the real-time FT checkpointing support is implemented by modifications and additions within a hypervisor. For example, if a virtualization-based real-time

checkpointing system was being implemented on Linux, it could be implemented via extensions and/or modifications to KVM/QEMU (a common Linux hypervisor). Various configurations of virtual machines (VM) may be used. Various components that support real-time checkpointing are shaded. In addition, various components or paths that are inactive are shown with a dotted line.

[0061] The system 70 includes Server 1 and Server 2. The servers may be any suitable computing devices. In one embodiment, the servers are commodity or custom servers running an OS which supports virtualization. Though the illustration shows a single hypervisor/VM on each server (a single real-time FT hypervisor/VM pair) various other configurations are possible. For example, the system 70 may include multiple hypervisor/VMs on each server and hence multiple real-time FT hypervisor/VM pairs, non-real-time hypervisor/VM pairs, and non-paired (unprotected) hypervisor/VMs. The system 70 includes a primary VM and a secondary VM associated with Server 1 and Server 2, respectively, as shown. The primary VM is in communication with the primary hypervisor. Similarly, the secondary VM is in communication with the secondary hypervisor. In general, the components of server 1 are mirrored in server 2 in terms of the functional blocks and software components shown relative to the hypervisor, checkpoint engine, storage support and network support features shown. In various embodiments, all of the foregoing are implemented using combinations of software modules and software programs.

[0062] Server 1 and Server 2 are connected or linked by a hardware interconnect. In one embodiment, the hardware interconnect is a physical medium such as a hardware bus or bridge through which checkpointing state is transferred between the two physical servers. The hardware interconnect can be a network connection, a non-transparent PCI bridge, or any other high-performance computer interconnect. In various embodiments, a checkpoint is transferred using the hardware interconnect during the occurrence of a low priority time slot being the active time slot on the TSB. In one embodiment, checkpoint data may be transferred on the hardware interconnect during all time slots (though transmission of the final state on the hardware interconnect does occur during the low priority time slot.) The "final stage" of a checkpoint is the most processor / system intensive operation and that other stages of the checkpoint do not impact real-time functionality. In various embodiments, one or more stages of a checkpoint, such as the final stage has a performance impact because the customer application or VM is not running at that time. As a result, the goal is to transfer the checkpoint data as quickly as possible which has associated performance costs.

[0063] As noted herein, such performance costs may negatively impact real-time performance. The connection of the hardware interconnect to various components of each computing device can vary depending on the type of hardware interconnect that is in use. For example,

if the interconnect is a non-transparent PCI bridge, it does not go through the network such as is shown at the bottom of FIG. 8. If the hardware interconnect is implemented using a network connection, it does connect to the network shown at the bottom of FIG. 8 and may connect from networking interfaces associated with the OS Networking Support. Selectively transferring checkpoint data that negatively impacts real-time performance during low priority transfers on the TSB avoids additional strain on a given server, which facilitates real-time performance.

[0064] Each of server 1 and server 2 include a checkpoint engine. The checkpoint engine and all of the components depicted within the hypervisor for each of server 1 and server 2 are additions to or modifications of each respective hypervisor. The checkpoint engine initiates periodic checkpoints as determined by the TSB timing. The purpose of the checkpoint process is to update the state of the secondary VM to match the state of the primary VM, so that if the hardware running the primary VM were to fail the secondary node can take over without an application outage.

[0065] The steps of the checkpoint process include one or more of the following:

- Stopping OS execution on both the primary and secondary VMs.
- Notifying the storage and network components that a checkpoint is starting.
- Obtaining lists of VM memory modified since the last checkpoint by both the primary and secondary VMs. Calculate the union of these two lists.
- Using the union of modified memory list, send the modified VM memory from the primary to secondary hypervisors, via the Hardware Interface.
- On the primary VM, capturing the state of the VM processors and devices. Send the processor and device state from the primary to secondary hypervisor via the Hardware Interface.
- On both the primary and secondary hypervisors informing the Network Support and Storage Support components that the checkpoint process has completed, and that the VM is about to resume.
- Resuming VM execution on both the primary and secondary nodes.

[0066] It is important that the system and time slots be properly sized to support real-time checkpointing. Timeslots may be sized according to the application for which the real-time fault tolerant system is being used as well as the cycle time specified by a user. During the checkpoint process, if either the primary or secondary checkpoint engine detects that the checkpoint duration will exceed the time allotted and hence possibly interfere with real-time deadline commitments, the checkpoint will be aborted, and the system reverted to a simplex state to ensure real-time deadlines are met. Reverted to simplex state refers to when the secondary and primary stop

checkpointing and the primary runs alone (and unprotected since it is no longer checkpointing). If this occurs, one server will re-partner to the other server to restore FT operation. Similarly, if either the primary or secondary checkpoint engine detects the need to perform a checkpoint at other than the timeslot specified by the TSB timing, the system will revert to a simplex state to ensure real-time deadlines are met. If this occurs, the system will re-partner itself to restore FT operation. Examples of items that can cause an unscheduled checkpoint are: an improperly sized system where buffers associated with memory tracking being too small to record all the memory changes that occur between checkpoints or storage caches too small to cache all the storage changes between checkpoints.

[0067] Within the checkpoint engine are two subcomponents, a Memory Tracker and a Misc Device State Replication component. The memory tracker utilizes processor hardware features to track memory modified by the VM so the checkpointing engine knows what memory needs to be copied at a checkpoint. In turn, the Misc Device State Replication component captures VM device state in the primary hypervisor and applies to the VM in the secondary hypervisor as part of a checkpoint.

[0068] The checkpoint engine in the secondary hypervisor of server 2 waits for the complete checkpoint data to be received, including the dirty page data and the Miscellaneous Device State data, and then applies that data and resumes VM operation in the context of the new checkpoint.

Storage Support Components

[0069] The storage support components interface with the hypervisor storage functionality in order to support checkpointing. The specific checkpointing related components are a caching component and a replication component. The caching component caches storage writes occurring between each checkpoint from both primary and secondary VMs. The primary and secondary VMs are non-deterministic when running. As a result, these VMS might not be writing the same data to storage. For example, if VMs are processing transactions in a different order because of non-determinism, the time stamps for the transactions written to storage on the primary could be different than the timestamps written to storage on the secondary. This storage interaction may be referenced as a view of the storage. The views are in the cache until the checkpoint, after which (in the non-failure case) the primary's data/view is written to disk and the secondary's data/view is discarded. Within a run interval, the VMs' view of storage is allowed to diverge between the VMs, as each VM operates independently and potentially issues different write requests. Following each checkpoint, the secondary VM's storage is made identical to the primary VM's storage. Caching writes instead of providing them to the OS to be committed to media provides certain advantages. For example, one advantage

is that the secondary cache can simply be discarded following a checkpoint. It would also be problematic to commit primary writes to the media before taking a checkpoint, in case of primary node failure and eventual recovery from the secondary; in that case the primary cache would likewise be discarded.

[0070] In addition, a replication component is also used to support various storage functionality. A storage component supports replicating storage between primary and secondary VMs. At the start of each run interval, VM storage is identical. This requires an initial full synchronization operation before operation begins. Then, working in conjunction with the caching component of the storage support component, storage modifications are replicated from primary to secondary hypervisors as each checkpoint is taken, and committed to the secondary storage media via the OS. Each of the storage components may be implemented as one or more software programs, methods or routines.

Network Support Components

[0071] The network support components interface with the hypervisor network functionality in order to support checkpointing and real-time operation. These components include a replicator, a local outgoing packet queue, a release coordinator, a partner outgoing packet queue and a release notifier. Each of these may be implemented as one or more software programs, methods or routines.

[0072] In the primary hypervisor, the replicator is operable to take outgoing network packets from the primary VM and sends the packet to the primary's local outgoing packet queue and the secondary's partner outgoing packet queue. In the secondary hypervisor, the replicator is inactive and serves simply as a pass-thru for packets from the secondary guest to the secondary's local outgoing packet queue. Along with the packets, the replicator passes information identifying which TSB time slot to which each packet is assigned. The local outgoing packet queue is operable to queue packets from the local VM instance for possible later release by the release coordinator. Along with the packets, information identifying which TSB time-slot each packet is assigned to is queued. The partner outgoing packet queue component is operable to queue packets from the partner VM instance for possible later release by the release coordinator. It is only active on the secondary hypervisor.

[0073] Two alternatives for releasing outgoing packets are discussed above. These include as follows:

- a) If the system can tolerate large numbers of duplicate packets, both the primary and secondary RT/FT can unconditionally release the queued outgoing packets.
- b) If the system cannot tolerate large numbers of duplicate packets, only the primary releases the queued outgoing packets in the non-failure case.

If alternative b is in use, the partner out going packet queue (only active on the secondary) is receiving notifications from the primary release notifier when packet(s) are released from the primary. The corresponding packet(s) are then removed from the partner outgoing packet queue.

[0074] The release coordinator is operable to release packets from the appropriate queue at the appropriate time to the OS and hence to the external network. In all cases, a packet is released during its assigned TSN time slot.

[0075] In non-failure cases, packets from the previous epoch are released from either:

- a) both the primary hypervisor's *local outgoing packet queue* and the secondary hypervisor's *partner outgoing packet queue* if alternative a from above is in use; or
- b) just the primary hypervisor's *local outgoing packet queue* if alternative b above is in use

If alternative b from above is in use, the primary's release coordinator notifies the primary's release notifier of each packet released. After a delay to ensure the packet has in fact been released, the release notifier passes the notification onto secondary's partner out going packet queue. If alternative b from above is in use and the primary hypervisor or VM fails, the secondary release coordinator releases the remaining packets in the secondary's partner outgoing packet queue.

[0076] The release notifier is only used in option b above and is operable to notify the secondary of any packets released by the primary. The main function of the release notifier is to add sufficient delay to notification to the secondary to ensure the packet has in fact been released by the primary. This is to avoid any dropped packets after a failover due to the former secondary thinking a packet has been released when it has not.

OS-Based Checkpointing

[0077] FIG. 9 is a block diagram of an OS-based for implementing real-time fault tolerance. OS-based checkpointing is described, which supports real-time operation using various components as depicted and discussed with regard to FIG. 9. A primary node / primary server running a primary OS is connected by a hardware interconnect to a secondary node / secondary server running a secondary OS. As shown, OS-based real-time FT checkpointing support is implemented by modifications and additions within or below the operating system. For example, a Linux OS-Based checkpointing solution would be largely implemented in a Linux storage driver, a Linux network driver, and a Linux checkpoint driver. The system 75 of FIG. 9 compares to the earlier discussed virtualization-based real-time FT checkpointing of FIG. 8 wherein support is implemented via hypervisor additions or modification. Various components that sup-

port real-time checkpointing are shaded. In addition, various components or paths that are inactive are shown with a dotted line.

[0078] Conceptually, many of the components in FIG. 9 perform the same function of the similarly named components in FIG. 8. There are differences that arise from a virtualized approach versus an operating system based-approach. Because OS-based checkpointing synchronizes systems at the machine level (vs the VM level), there are more stringent constraints on the actual hardware for OS-based RT/FT. The Primary Node/Server and Secondary Node/Server are commodity or custom servers that satisfy one or more of the following requirements:

- The servers are configured with an identical hardware configuration
- The servers contain hardware supported by the checkpointing network driver, the checkpointing storage driver, and the checkpoint shim for the target OS.

Also, unlike virtualization-based checkpointing, wherein multiple VMs can be checkpointed and both physical servers can hence simultaneously host primary VMs and secondary VMs, OS-based checkpointing has a primary server and a secondary server.

[0079] The server OS is protected by and contains the checkpointing technology. This server OS can be a standard OS (Windows, Linux, and/or VMware) but various OS specific driver additions and modifications are needed to support the checkpointing functionality.

[0080] In one embodiment, the hardware interconnect is a physical medium such as a hardware bus or bridge through which checkpointing state is transferred between the two physical servers. The hardware interconnect can be a network connection, a non-transparent PCI bridge, or any other high-performance computer interconnect. In various embodiments, the hardware interconnect is separate from the TSB and does not connect thereto. Software support for the hardware interconnect is within the checkpoint engine driver. Native OS support for the hardware interface cannot be used because the checkpointing operations need to be independent of the checkpointing OS (for example, checkpoint activity needs to continue even if the host OS is generally paused for a checkpoint). This means for example, that if TCP/IP is used over a NIC to transfer the checkpoint driver, a TCP/IP stack distinct from the host OSs TCP/IP stack is implemented in the checkpoint driver.

[0081] Referring to FIG. 9, the checkpoint engine initiates periodic checkpoints as determined by the TSB timing. The purpose of the checkpoint process is to update the state of the secondary node to match the state of the primary node, so that if the primary node were to fail, the secondary node can take over without an application outage.

[0082] The steps of the checkpoint process are:

- Informing Network and Storage drivers that a checkpoint is imminent. The checkpoint pauses here until storage driver performs callback indicating outstanding storage accesses have been completed.
- Stopping OS execution on both the primary and secondary nodes.
- Obtaining lists of memory modified since the last checkpoint on both the primary and secondary nodes. Calculate the union of these two lists.
- Using the union of modified memory list, send the modified memory from the primary to secondary nodes, via the Hardware Interface.
- On the primary node, capturing the state of the processors and devices.
- Sending the processor and device state from the primary to secondary nodes via the Hardware Interface.
- On both the primary and secondary nodes informing the Network and Storage drivers that the checkpoint process has completed, and that the OS is about to resume.
- Resuming OS execution on both the primary and secondary nodes.

[0083] In various implementations of OS-based RT/FT, the system and time slots should be properly sized to support real-time checkpointing. During the checkpoint process, if either the primary or secondary checkpoint driver detects that the checkpoint duration will exceed the time allotted on the TSB for low priority communication (and hence possibly interfere with real-time deadline commitments), the checkpoint will be aborted and the system reverted to a simplex state to ensure real-time deadlines are met. If this occurs, the system will re-partner itself to restore FT operation.

[0084] Similarly, if either the primary or secondary checkpoint driver detects the need to perform a checkpoint at other than the timeslot specified by the TSB timing, the system will revert to a simplex state to ensure real-time deadlines are met. If this occurs, the system will re-partner itself to restore FT operation. Examples of items that can cause an unscheduled checkpoint are buffers associated with memory tracking being too small to record all the memory changes that occur between checkpoints, and storage caches too small to cache all the storage changes between checkpoints.

[0085] The Checkpoint Shim is transparent to the OS and is a container for the Memory Tracker and Misc Device State Replication. This checkpoint shim component tracks memory modified by server OS so the checkpointing engine knows what memory needs to be copied at a checkpoint. In one embodiment the checkpoint shim utilizes Intel's PML (Page Modification Logging) hardware feature to track memory. In various embodiments, other hardware specific features such as motherboard or processors or other chipset specific instructions and features may be used to track memory. VT-x hardware virtualization extensions are enabled when using PML under certain scenarios. If VT-x is enabled, hardware reserves a

contiguous area of physical memory for virtual machine monitor (VMM) code to operate. In a typical virtual machine environment VMM code is invoked when a virtualized guest operating system executes an instruction that needs special handling, such as interrupts, memory faults, IO access, certain privileged instructions, and indications that the PML generated list that includes modified memory locations needs to be serviced.

[0086] The Memory Tracker installs a micro-VMM shim (a minimal hypervisor) having a main function of handling PML FULL vmexit events and maintain dirty-page buffers in a reserved memory region. The VMM also processes a number of other vmexit events from the OS, mainly those vmexits which cannot be disabled in hardware while using PML. These vmexit events are emulated within the micro-VMM so that the server will continue to operate in its normal configuration while the VMM is present.

[0087] When the system workload has generated enough dirtied pages to fill up the hardware PML log, a PML FULL vmexit will occur, the new set of pages will be tracked, and workload processing will continue. This enables the Memory Tracker to generate a list of physical memory pages which have been modified on each Server since the previous checkpoint. The memory tracking embodiment tracks memory modifications at the page level; other embodiments may track memory at a different granularity.

[0088] Misc Device State Replication is a checkpoint shim component. Specifically, this checkpoint shim component captures device state not captured by OS resident components from one node/server and applies it to the other node/server as part of a checkpoint.

[0089] The checkpoint engine on the primary server issues a request to the VMM, to read or extract some device and processor state, such as timestamp counters, processor registers, stack, and other data, and to record that Misc Device State data in a reserved memory region. This data is transferred to the secondary server.

[0090] The checkpoint engine on the secondary side requests the VMM to wait for the complete checkpoint data to be received, including the dirty page data and the Miscellaneous Device State data, and then the VMM applies that data and resumes to the OS in the context of the new checkpoint.

Storage Driver and Sub-Components

[0091] The storage driver is an OS storage hardware driver designed to support checkpointing. Specific components in the support of checkpointing are a caching component and a replication component, described immediately below.

[0092] The caching components of the OS storage driver caches storage writes occurring between each checkpoint, on both primary and secondary nodes. Within a checkpoint run interval, the OS's view of storage is allowed to diverge between the nodes, as each node operates independently and potentially issues different

write requests. Following each checkpoint, the secondary storage is made identical to the primary. Caching writes instead of committing them to the media provides various advantages. For example, the secondary cache can simply be discarded following a checkpoint. It would also be problematic to commit primary writes to the media before taking a checkpoint, in case of primary node failure and eventual recovery from the secondary; in that case the primary cache would likewise be discarded.

[0093] The replication component of the OS storage driver supports replicating storage between primary and secondary nodes. At the start of each run interval, storage is identical. This requires an initial full synchronization operation before operation begins. Then, working in conjunction with the caching component of the driver, storage modifications are replicated from primary to secondary nodes as each checkpoint is taken, and committed to the secondary storage media.

Network Driver and Sub-Components

[0094] The network driver is an OS network hardware driver designed to support real-time FT checkpointing. The network driver's components include a replicator, a local outgoing packet queue, a release coordinator, a partner outgoing packet queue, and a release notifier. Each of these may be implemented as one or more software programs, methods or routines.

[0095] On the primary node the replicator component takes outgoing network packets from the OSs network stack and sends the packet to the primary's local outgoing packet queue and the secondary's partner outgoing packet queue. On the secondary node, the replicator component is inactive and serves simply as a pass-thru for packets from the secondary OSs network stack to the secondary's local outgoing packet queue. Along with the packets, the replicator passes information identifying which TSB time slot each packet is assigned to.

[0096] The local outgoing packet queue component queues packets from the local OS network stack for possible later release by the release coordinator. Along with the packets, information identifying which TSB time-slot each packet is assigned to is queued. The partner outgoing packet queue component queues packets from the partner OS network stack for possible later release by the release coordinator. It is only active on the secondary node.

[0097] Two alternatives for releasing outgoing packets are discussed above. These also apply to the OS-based RT/FT systems. These include as follows:

- a) If the system can tolerate large numbers of duplicate packets, both the primary and secondary RT/FT can unconditionally release the queued outgoing packets.
- b) If the system cannot tolerate large numbers of duplicate packets, only the primary releases the queued outgoing packets in the non-failure case.

[0098] If alternative b is in use, the partner outgoing packet queue (only active on the secondary) is receiving notifications from the primary release notifier when packet(s) are released from the primary. The corresponding packet(s) are then removed from the partner outgoing packet queue.

[0099] The release coordinator releases packets from the appropriate queue at the appropriate time to the OS and hence to the external network. In all cases, a packet is released only during its assigned TSB time slot.

[0100] In non-failure cases, packets from the previous epoch are released from either both the primary node's local outgoing packet queue and the secondary node's partner outgoing packet queue, if alternative a from above is in use just the primary hypervisor's local outgoing packet queue if alternative b above is in use

[0101] If alternative b from above is in use, the primary's release coordinator notifies the primary's release notifier of each packet released. After a delay to ensure the packet has in fact been released, the release notifier passes the notification onto secondary's partner outgoing packet queue.

[0102] If alternative b from above is in use and the primary hypervisor or VM fails, the secondary Release Coordinator releases the remaining packets in the secondary's partner outgoing packet queue.

[0103] The release notifier is only used in option b above and notifies the secondary of any packets released by the primary. The main function of the release notifier is to add sufficient delay to notification to the secondary to ensure the packet has in fact been released by the primary. This is to avoid any dropped packets after a failover due to the former secondary thinking a packet has been released when it has not.

Variations and Alternative Embodiments

[0104] Standard active-active checkpointing (e.g. CO-LO) has two performance advantages relative to the real-time checkpointing described in the previous sections:

- Network packets are not gated, thereby significantly reducing network latency and thereby improving performance
- Checkpoints are generally only called if there is a sufficient divergence in an outgoing packet produced by both nodes, reducing the number of checkpoints for most workloads.

[0105] It is not possible to implement a real-time fault tolerant solution using active-active checkpointing because the two nodes (OS-based checkpointing) or 2 VMs (virtualization-based checkpointing) may produce a sufficiently divergent packet at any time and hence require a checkpoint at any time, thereby causing a real-time deadline to be missed. Notwithstanding the foregoing, it is possible to relax the two performance advantage noted above at the cost of some tradeoffs, creating three addi-

tional real-time FT checkpointing alternatives. These three alternatives or variants include the following embodiments:

1) Variant 1 - reduce checkpoints

- Skip checkpoints if not needed

2) Variant 2 - reduce network gating

- Do not gate packets unless there is a miscompare

3) Variant 3 (variants 1 and 2 combined)

- Do not gate packets unless there is a miscompare
- Skip checkpoints if not needed

[0106] The three variants listed above require the addition of logic to compare outgoing packets between the primary and secondary nodes. Specifically, for virtualization-based real-time checkpointing shown in FIG. 10 logic to compare outgoing packets from the secondary node to those from the primary node is added to Network Support component. Specifically, packet comparison logic is added. This packet comparison logic only needs to be active on one node though both nodes need to know the result of the comparison. Similarly, FIG. 11 shows a diagram of the OS-based real-time checkpointing network driver with the addition of packet comparison logic added. In general, references to partner refers to the other computing device or virtual machine / hypervisor as applicable.

Variant 1 - Reduce Checkpoints

[0107] If outgoing packets produced by both the primary and secondary are effectively equivalent during a cycle there is no need to perform a checkpoint. This has the advantage that compute cycles and run time that would otherwise be spent on the checkpoint are available for the application. It does have disadvantages that include one or more of the following:

- 1) Logic is added to the design to compare outgoing packets from the primary and secondary.
- 2) When a checkpoint is needed, it will be on average longer than if a checkpoint was taken every cycle. The minimum designed possible cycle time therefore is increased to allow for these longer checkpoints.
- 3) Due to reasons outlined immediately below, packets are gated for up to a little over two cycles if a miscompare occurs too late in a cycle to call a checkpoint. As described later, this breaks the regular cyclic behavior of the base implementation and can also further increase the minimum cycle time.

[0108] FIG. 12 A illustrates the need for the packet gating if a miscompare occurs too late in a cycle to call a checkpoint. As shown in cycle 3 of FIG. 12B, if packets between the primary and secondary miscompare when it's too late to call a checkpoint the two nodes enter the next cycle in an inconsistent state. As a result, inconsistent state is released to the outside world if the secondary has to take over for the primary (due to a failure of the primary) in the following cycle.

[0109] One solution to this problem is to start gating packets once the miscompare occurs, take a checkpoint the following cycle, and continue to gate the packets until after the checkpoint. This prevents any detectable inconsistent state from being transmitted to the outside world thereby ensuring the secondary is in a sufficiently consistent state to take over for the primary. If the primary fails, it's the secondary's packets which are released. This adds at least two cycles of latency to the packets, one cycle for the delay just mentioned, and another cycle because this variant, like the base solution, always adds a cycle of delay. This will cause a cycle with no packets transmitted and a subsequent cycle with two cycles worth of packets being transmitted. This breaks cyclic behavior and could impact the minimum possible cycle time if transmitting twice the number of packets in a cycle is on the critical timing path.

[0110] While Variant 1 provides additional compute cycles, the additional compute cycles are discretionary compute cycles. The system is sized so that it can accomplish the minimum required work assuming a checkpoint in every cycle. In practice though, assuming relatively short cycle times, the vast majority of cycles should not require a checkpoint.

Variant 2 - Reduce Network Gating

[0111] In one embodiment, a real-time checkpointing solution gates all outgoing packets every run interval. This can be relaxed to only start gating packets once a miscompare occurs. This has the advantage of reducing average network latency. In various embodiments, as part of the implementation of this variant, logic, such as hardware or software modules is added to the applicable systems and methods to compare outgoing packets from the primary and secondary. Network traffic is no longer strictly cyclic as illustrated in FIG. 12B. Specifically, referring to FIG. 12B, no packets will be delivered during time slot 3 in cycle 2, whereas in timeslot 3 of cycle 3, packets from both cycles 2 and 3 will be released. If network bandwidth was critical path in determining minimum cycle time, then the cycle time is increased because two cycles worth of packets may have to be released in a single time slot.

[0112] It should also be noted that it is only average network latency which is reduced, worst case latency remains the same or is worse. This occurs because the very first packet may miscompare and cycle times are possibly longer due to network bandwidth being a critical

path in determining minimum cycle time.

Variant 3 - Reduce Network Gating and Checkpoints

[0113] Variants 1 and 2 may be combined to produce a solution with lower average network latency and more discretionary compute cycles, though it is subject to the disadvantages of variant 1 and 2 outlined above.

Improvements of Real-time Fault Tolerant System Relative to Legacy Checkpointing

[0114] The real-time fault tolerant systems and methods discussed above improve upon and avoid features of various legacy checkpointing approaches. Active-passive checkpointing solutions have an active VM or OS running the target workload, and a passive (standby) VM or OS which is effectively idle and is receiving periodic snapshots of the active VM or OS's state in order to pick up running the target workload should the prior active VM or OS fail.

[0115] Active-passive checkpointing solutions synchronize VM or OS state at high frequencies, buffering (gating) the output network packets during each epoch until checkpoint completes (state capture and transfer) to ensure external clients are supplied functionally consistent network packets across a failover. Further, active-active checkpointing improves performance relative to active-passive checkpointing for most workloads by relaxing the frequency of checkpoints based on output similarity and by removing the need to gate the output network packets during each epoch.

[0116] In active-active checkpointing systems, both VM and/or OS instances are running, outgoing network packets are compared between the instances and immediately released if they are sufficiently similar. The redundant VMs or OSs may diverge as long as they generate effectively identical responses to client requests. If corresponding outgoing network packets are no longer effectively identical, a checkpoint is called to re-synchronize the VMs or OSs.

[0117] Active-passive checkpointing is not suitable for real-time operation for multiple reasons. Similarly, active-active checkpointing is not suitable for real-time operation for multiple reasons. Several of these reason are discussed below and provide evidence as to why legacy active-passive and active-active checkpointing differ from various RT/FT embodiments of the presence disclosure because such RT/FT embodiments avoid one, or more or all of the following modes of operation.

[0118] In an active-passive system, checkpoint timing is independent of TSB timing and can cause real-time deadlines to be missed. If a failure occurs, the system reverts back to the last checkpoint, i.e. up to an epoch of processing time is lost, possibly causing a real-time deadline to be missed. In addition, active-passive checkpointing can drop outgoing packets upon a failure. This breaks real-time commitments and hence is not accept-

able for most real-time applications. Non-real-time solutions rely on higher-level network protocols (e.g. TCP) to compensate for the dropped packets. With regard to active-active checkpointing, checkpoints occur whenever a packet miscompare occurs independent of TSB timing and hence can cause real-time deadlines to be missed. A given RT/FT embodiment operates in a manner that avoids the active-passive and active-active checkpointing issues and design requirements discussed herein.

[0119] A fault tolerant computer system frequently includes two processor subsystems, generally referred to as the primary and secondary nodes. These nodes either operate substantially synchronously executing the same program in parallel (active-active) or operating with one system, the primary node, actively doing useful computations with the secondary node standing by (active-standby) waiting for a failure of the primary system. If the primary system fails the secondary takes over as the new primary and continues the work of the failing primary, allowing it to be replaced.

General and Specific Computer-based Implementations and Embodiments

[0120] Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "checkpointing" or "calculating" or "delaying" or "comparing", "generating" or "determining" or "forwarding" or "deferring" "committing" or "interrupting" or "handling" or "receiving" or "buffering" or "allocating" or "displaying" or "flagging" or Boolean logic or other set related operations or the like, refer to the action and processes of a computer system, or electronic device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's or electronic devices' registers and memories into other data similarly represented as physical quantities within electronic memories or registers or other such information storage, transmission or display devices.

[0121] The algorithms presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems is apparent from the description above. In addition, the disclosure is not described with reference to any particular programming language, and various embodiments may thus be implemented using a variety of programming languages.

[0122] A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the disclosure. For example, various forms of the flows shown above may be used, with steps reordered, added, or removed. Accordingly, other imple-

mentations are within the scope of the following claims.

[0123] The examples presented herein are intended to illustrate potential and specific implementations of the disclosure. The examples are intended primarily for purposes of illustration of the disclosure for those skilled in the art. No particular aspect or aspects of the examples are necessarily intended to limit the scope of the disclosure.

[0124] The figures and descriptions of the disclosure have been simplified to illustrate elements that are relevant for a clear understanding of the disclosure, while eliminating, for purposes of clarity, other elements. Those of ordinary skill in the art may recognize, however, that these sorts of focused discussions would not facilitate a better understanding of the disclosure, and therefore, a more detailed description of such elements is not provided herein.

[0125] The processes associated with the present embodiments may be executed by programmable equipment, such as computers. Software or other sets of instructions that may be employed to cause programmable equipment to execute the processes may be stored in any storage device, such as, for example, a computer system (non-volatile) memory, an optical disk, magnetic tape, or magnetic disk. Furthermore, some of the processes may be programmed when the computer system is manufactured or via a computer-readable memory medium.

[0126] It can also be appreciated that certain process aspects described herein may be performed using instructions stored on a computer-readable memory medium or media that direct a computer or computer system to perform process steps. A computer-readable medium may include, for example, memory devices such as diskettes, compact discs of both read-only and read/write varieties, optical disk drives, and hard disk drives. A computer-readable medium may also include memory storage that may be physical, virtual, permanent, temporary, semi-permanent and/or semi-temporary.

[0127] Computer systems and computer-based devices disclosed herein may include memory for storing certain software applications used in obtaining, processing, and communicating information. It can be appreciated that such memory may be internal or external with respect to operation of the disclosed embodiments. The memory may also include any means for storing software, including a hard disk, an optical disk, floppy disk, ROM (read only memory), RAM (random access memory), PROM (programmable ROM), EEPROM (electrically erasable PROM) and/or other computer-readable memory media. In various embodiments, a "host," "engine," "loader," "filter," "platform," or "component" may include various computers or computer systems, or may include a reasonable combination of software, firmware, and/or hardware.

[0128] In various embodiments of the disclosure, a single component may be replaced by multiple components, and multiple components may be replaced by a single component, to perform a given function or functions. Ex-

cept where such substitution would not be operative to practice embodiments of the disclosure, such substitution is within the scope of the disclosure. Any of the servers, for example, may be replaced by a "server farm" or other grouping of networked servers (e.g., a group of server blades) that are located and configured for cooperative functions. It can be appreciated that a server farm may serve to distribute workload between/among individual components of the farm and may expedite computing processes by harnessing the collective and cooperative power of multiple servers. Such server farms may employ load-balancing software that accomplishes tasks such as, for example, tracking demand for processing power from different machines, prioritizing and scheduling tasks based on network demand, and/or providing backup contingency in the event of component failure or reduction in operability.

[0129] In general, it may be apparent to one of ordinary skill in the art that various embodiments described herein, or components or parts thereof, may be implemented in many different embodiments of software, firmware, and/or hardware, or modules thereof. The software code or specialized control hardware used to implement some of the present embodiments is not limiting of the disclosure. Programming languages for computer software and other computer-implemented instructions may be translated into machine language by a compiler or an assembler before execution and/or may be translated directly at run time by an interpreter.

[0130] Examples of assembly languages include ARM, MIPS, and x86; examples of high level languages include Ada, BASIC, C, C++, C#, COBOL, Fortran, Java, Lisp, Pascal, Object Pascal; and examples of scripting languages include Bourne script, JavaScript, Python, Ruby, PHP, and Perl. Various embodiments may be employed in a Lotus Notes environment, for example. Such software may be stored on any type of suitable computer-readable medium or media such as, for example, a magnetic or optical storage medium. Thus, the operation and behavior of the embodiments are described without specific reference to the actual software code or specialized hardware components. The absence of such specific references is feasible because it is clearly understood that artisans of ordinary skill would be able to design software and control hardware to implement the embodiments of the disclosure based on the description herein with only a reasonable effort and without undue experimentation.

[0131] Various embodiments of the systems and methods described herein may employ one or more electronic computer networks to promote communication among different components, transfer data, or to share resources and information. Such computer networks can be classified according to the hardware and software technology that is used to interconnect the devices in the network.

[0132] The computer network may be characterized based on functional relationships among the elements or components of the network, such as active networking, client-server, or peer-to-peer functional architecture. The

computer network may be classified according to network topology, such as bus network, star network, ring network, mesh network, star-bus network, or hierarchical topology network, for example. The computer network may also be classified based on the method employed for data communication, such as digital and analog networks.

[0133] Embodiments of the methods, systems, and tools described herein may employ internetworking for connecting two or more distinct electronic computer networks or network segments through a common routing technology. The type of internetwork employed may depend on administration and/or participation in the internetwork. Non-limiting examples of internetworks include intranet, extranet, and Internet. Intranets and extranets may or may not have connections to the Internet. If connected to the Internet, the intranet or extranet may be protected with appropriate authentication technology or other security measures. As applied herein, an intranet can be a group of networks which employ Internet Protocol, web browsers and/or file transfer applications, under common control by an administrative entity. Such an administrative entity could restrict access to the intranet to only authorized users, for example, or another internal network of an organization or commercial entity.

[0134] Unless otherwise indicated, all numbers expressing lengths, widths, depths, or other dimensions and so forth used in the specification and claims are to be understood in all instances as indicating both the exact values as shown and as being modified by the term "about." As used herein, the term "about" refers to a $\pm 10\%$ variation from the nominal value. Accordingly, unless indicated to the contrary, the numerical parameters set forth in the specification and attached claims are approximations that may vary depending upon the desired properties sought to be obtained. At the very least, and not as an attempt to limit the application of the doctrine of equivalents to the scope of the claims, each numerical parameter should at least be construed in light of the number of reported significant digits and by applying ordinary rounding techniques. Any specific value may vary by 20%.

[0135] In various embodiments, one or more techniques, methods, systems, subsystems, and combinations of the foregoing can be implemented using virtualized operating systems, container-based operating systems, virtualized environments, and various emulators, such as VMware, QEMU, KVM, and others.

[0136] It will be appreciated by those skilled in the art that various modifications and changes may be made without departing from the scope of the described technology. Such modifications and changes are intended to fall within the scope of the embodiments that are described. It will also be appreciated by those of skill in the art that features included in one embodiment are interchangeable with other embodiments; and that one or more features from a depicted embodiment can be included with other depicted embodiments in any combi-

nation. For example, any of the various components described herein and/or depicted in the figures may be combined, interchanged, or excluded from other embodiments.

Claims

1. A real-time fault tolerant computing system (7b) comprising:

a first computing device (8a) comprising

one or more memory devices,
one or more processors,
a first network interface operable to receive device data and transmit output data over a time-slot-based bus, wherein the output data is generated from processing device data, and wherein the bus is operable to receive data according to time slots of two or more priority levels, and

a first real-time checkpoint engine;

a second computing device (8b) comprising

one or more memory devices,
one or more processors,
a second network interface operable to receive device data over the time slot-based bus, and
a second real-time checkpoint engine; and

a hardware interconnect operable to permit data exchange between the first computing device and the second computing device, wherein the first computing device (8a) and the second computing device (8b) simultaneously process received device data, and wherein checkpoints are generated by the first or the second real-time checkpoint engines during lower-priority communication time slots allocated on the time slot-based bus to avoid interfering with any real-time communications to or from the first and second computing devices, wherein the lower-priority communication time slots comprise non-real-time communication time slots.

2. The real-time fault tolerant computing system (7b) of claim 1 wherein if the first computing device (8a) experiences a failure the second computing device continues processing device data to generate the output data and transmit the output data using the time slot-based bus.

3. The real-time fault tolerant computing system (7b) of claim 2 wherein checkpoint operations are always

postponed until the occurrence of a lower-priority communication slot, even upon the first computing device (8a) and the second computing device (8b) producing divergent output data prior to the lower-priority communication slot occurring.

4. The real-time fault tolerant computing system (7b) of claim 3 wherein the first computing device (8a) and the second computing device (8b) both delay release of output data to one or more devices in communication with the time-slot based bus until after a subsequent checkpoint to avoid release of divergent output data should a failure occur.
5. The real-time fault tolerant computing system (7b) of claim 2, wherein the first computing device (8a) generates checkpoints on a per cycle basis when output data comprising low priority data is transmitted using a low priority time slot.
6. The real-time fault tolerant computing system (7b) of claim 2, wherein the device is a sensor and the output data comprises sensor data, the sensor data is operable as an input for a control system.
7. The real-time fault tolerant computing system (7b) of claim 1, wherein the first computing device (8a) and the second computing device (8b) are both actively processing device data and producing output data so that if the first computing device (8a) experiences one or more failure modes a time delay for second processing device (8b) to continue processing is zero.
8. The real-time fault tolerant computing system (7b) of claim 2, wherein duration of one or more time slots is increased by a detection period, wherein the detection period ranges from about 1 microsecond to about 100 microseconds.
9. The real-time fault tolerant computing system (7b) of claim 2, wherein the first computing device (8a) and the second computing device (8b) both delay release of output data to devices in communication with the time-slot based bus until after the next checkpoint to avoid release of inconsistent output data, and optionally wherein if the first computing device (8a) is operating normally after a checkpoint, the first computing device (8a) releases any delayed output data.
10. The real-time fault tolerant computing system (7b) of claim 2, wherein the first computing device (8a) transmits all of its output data to the second computing device (8b) to avoid such output data being dropped upon occurrence of a failure of the first computing device (8a), and optionally wherein if the first computing device (8a) fails, the second computing

device (8b) takes over release of the first computing device's output data that was generated during the last checkpoint.

11. The real-time fault tolerant computing system (7b) of claim 2, wherein only the first computing device (8a) handles release of output data when no failures have occurred and in parallel shares such output data with the second computing device (8b).
12. The real-time fault tolerant computing system (7b) of claim 1, wherein the device data from one or more devices are transmitted to the first computing device (8a) and the second computing device (8b) over the time slot-based bus.
13. The real-time fault tolerant computing system (7b) of claim 2, wherein the first computing device (8a) and the second computing device (8b) receive incoming data directly through a broadcast or a multicast.
14. The real-time fault tolerant computing system (7b) of claim 2, wherein each of the first computing device (8a) and the second computing device (8b) comprise a respective non-virtual operating system (OS) comprising a respective network driver, checkpoint shim, and storage driver stored in the one or more memory devices and executing on each of the one or more processors, and optionally wherein the network driver comprises a replicator in communication with a release coordinator and a release notifier.
15. The real-time fault tolerant computing system (7b) of claim 3, wherein one or more networking programs stored in the one or more memory devices of the first and second computing devices delay and coordinate transmission of output data using one or more of the hardware interconnect, the first network interface and the second network interface.
16. The real-time fault tolerant computing system (7b) of claim 2, wherein each of the first computing device (8a) and the second computing device (8b) comprise a respective virtual operating system (OS) comprising one or more respective network support applications, a checkpoint engine, and one or more storage support applications stored in the one or more memory devices and executing with respective hypervisors.
17. The system of claim 1, wherein checkpoints are transmitted using the hardware interconnect.
18. A method of controlling an operational system in real-time comprising receiving device data, at a primary node and a secondary node, from one or more devices in electrical communication with a time slot-

based bus, wherein the bus is operable to receive data according to time slots of two or more priority levels; processing the received device data simultaneously at the primary node and the secondary node; during a first time slot, generating a checkpoint at the primary node and transmitting the checkpoint data to the secondary node, wherein the first time slot is a lower priority time slot, wherein the lower-priority communication time slot comprises a non-real-time time slot; generating an output data at the primary node and the secondary node in response to processing the received device data; and transmitting, from the primary node, the output data over the bus using a second time slot, wherein the second time slot is a higher priority time slot, wherein both the first node and the second node are in an active mode.

19. The method of claim 18 further comprising gating network data using a respective network software module executing on both the primary and the secondary node such that output data is released from the node whose state has carried forward to current checkpoint.
20. The method of claim 19 further comprising synchronizing transmission of output data with high priority time slots, wherein the output data is received by the operational system.
21. The method of claim 18 or claim 19 further comprising transmitting data during a non-checkpoint portion of a lower priority time slot.
22. The method of claim 18 further comprising transmitting, using the primary node, all of its respective outgoing data to the secondary node to avoid such outgoing data from being dropped upon occurrence of a failure of the primary node.

Patentansprüche

1. Fehlertolerantes Echtzeitrechensystem (7b), das Folgendes umfasst:

ein erstes Rechengerät (8a), das Folgendes umfasst:

ein oder mehrere Speichergeräte,
einen oder mehrere Prozessoren,
eine erste Netzwerkschnittstelle mit der Aufgabe, Gerätedaten zu empfangen und Ausgabedaten über einen zeitschlitzbasierten Bus zu übertragen, wobei die Ausgabedaten aus der Verarbeitung von Gerätedaten erzeugt werden und wobei der Bus die Aufgabe hat, Daten gemäß Zeitschlitzen

von zwei oder mehr Prioritätsstufen zu empfangen, und

eine erste Echtzeit-Checkpoint-Engine;
eine zweites Rechengerät (8b), das Folgendes umfasst:

ein oder mehrere Speichergeräte,
einen oder mehrere Prozessoren,
eine zweite Netzwerkschnittstelle mit der Aufgabe, Gerätedaten über den zeitschlitzbasierten Bus zu empfangen, und
eine zweite Echtzeit-Checkpoint-Engine;
und

eine Hardware-Verbindung mit der Aufgabe, einen Datenaustausch zwischen dem ersten Rechengerät und dem zweiten Rechengerät zuzulassen,

wobei das erste Rechengerät (8a) und das zweite Rechengerät (8b) gleichzeitig empfangene Gerätedaten verarbeiten, und
wobei Checkpoints von der ersten oder der zweiten Echtzeit-Checkpoint-Engine während Kommunikationszeitschlitzen niedrigerer Priorität erzeugt werden, die auf dem zeitschlitzbasierten Bus zugewiesen sind, um eine Störung jeglicher Echtzeitkommunikationen zu oder von dem ersten und zweiten Rechengerät zu vermeiden, wobei die Kommunikationszeitschlitze niedrigerer Priorität Nicht-Echtzeit-Kommunikationszeitschlitze umfassen.

2. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 1, wobei, wenn das erste Rechengerät (8a) einen Ausfall erfährt, das zweite Rechengerät die Verarbeitung von Gerätedaten fortsetzt, um die Ausgabedaten zu erzeugen und die Ausgabedaten über den zeitschlitzbasierten Bus zu übertragen.

3. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei Checkpoint-Operationen immer bis zum Auftreten eines Kommunikationsschlitzes niedrigerer Priorität verschoben werden, selbst wenn das erste Rechengerät (8a) und das zweite Rechengerät (8b) vor dem Auftreten des Kommunikationsschlitzes niedrigerer Priorität divergierende Ausgabedaten produzieren.

4. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 3, wobei sowohl das erste Rechengerät (8a) als auch das zweite Rechengerät (8b) die Freigabe von Ausgabedaten an ein oder mehrere Geräte in Kommunikation mit dem zeitschlitzbasierten Bus bis nach einem nachfolgenden Checkpoint verzögern, um die Freigabe divergierender Ausgabedaten zu vermeiden, falls ein Ausfall auftritt.

5. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei das erste Rechengerät (8a) Checkpoints auf einer Pro-Zyklus-Basis erzeugt, wenn Daten niedriger Priorität umfassende Ausgabedaten unter Verwendung eines Zeitschlitzes niedriger Priorität übertragen werden. 5
6. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei das Gerät ein Sensor ist und die Ausgabedaten Sensordaten umfassen, wobei die Sensordaten als ein Eingang für ein Steuersystem verwendet werden können. 10
7. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 1, wobei sowohl das erste Rechengerät (8a) als auch das zweite Rechengerät (8b) aktiv Gerätedaten verarbeiten und Ausgabedaten produzieren, so dass, wenn das erste Rechengerät (8a) einen oder mehrere Fehlermodi erfährt, eine Zeitverzögerung für das zweite Verarbeitungsgesamt (8b) zum Fortsetzen der Verarbeitung Null ist. 15 20
8. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei die Dauer eines oder mehrerer Zeitschlitzes um eine Erfassungsperiode erhöht wird, wobei die Erfassungsperiode im Bereich von etwa 1 Mikrosekunde bis etwa 100 Mikrosekunden reicht. 25
9. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei sowohl das erste Rechengerät (8a) als auch das zweite Rechengerät (8b) die Freigabe von Ausgabedaten an Geräte in Kommunikation mit dem zeitschlitzbasierten Bus bis nach dem nächsten Checkpoint verzögern, um die Freigabe inkonsistenter Ausgabedaten zu vermeiden, und optional wobei, wenn das erste Rechengerät (8a) nach einem Checkpoint normal arbeitet, das erste Rechengerät (8a) verzögerte Ausgabedaten freigibt. 30 35
10. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei das erste Rechengerät (8a) alle seine Ausgabedaten an das zweite Rechengerät (8b) überträgt, um zu vermeiden, dass solche Ausgabedaten beim Auftreten eines Fehlers des ersten Rechengeräts (8a) verworfen werden, und wobei optional, wenn das erste Rechengerät (8a) ausfällt, das zweite Rechengerät (8b) die Freigabe der Ausgabedaten des ersten Rechengeräts übernimmt, die während des letzten Checkpoints erzeugt wurden. 40 45
11. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei nur das erste Rechengerät (8a) die Freigabe von Ausgabedaten übernimmt, wenn keine Ausfälle aufgetreten sind, und solche Ausgabedaten parallel dazu mit dem zweiten Rechengerät (8b) teilt. 50 55
12. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 1, wobei die Gerätedaten von einem oder mehreren Geräten über den zeitschlitzbasierten Bus zum ersten Rechengerät (8a) und zum zweiten Rechengerät (8b) übertragen werden.
13. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei das erste Rechengerät (8a) und das zweite Rechengerät (8b) eingehende Daten direkt über einen Broadcast oder einen Multicast empfangen.
14. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei sowohl das erste Rechengerät (8a) als auch das zweite Rechengerät (8b) ein jeweiliges nicht-virtuelles Betriebssystem (OS) umfassen, das einen jeweiligen Netzwerktreiber, ein Checkpoint-Shim und einen Speichertreiber umfasst, die in den ein oder mehreren Speichergeräten gespeichert sind und auf jedem der ein oder mehreren Prozessoren ausgeführt werden, und wobei der Netzwerktreiber optional einen Replikator in Kommunikation mit einem Freigabekoordinator und einem Freigabebenachrichtiger umfasst.
15. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 3, wobei ein oder mehrere Netzwerkprogramme, die in den ein oder mehreren Speichergeräten des ersten und zweiten Rechengeräts gespeichert sind, die Übertragung von Ausgabedaten über eines oder mehrere von der Hardware-Verbindung, der ersten Netzwerkschnittstelle und der zweiten Netzwerkschnittstelle verzögern und koordinieren.
16. Fehlertolerantes Echtzeitrechensystem (7b) nach Anspruch 2, wobei sowohl das erste Rechengerät (8a) als auch das zweite Rechengerät (8b) ein jeweiliges virtuelles Betriebssystem (OS) umfassen, das eine oder mehrere jeweilige Netzwerkunterstützungsanwendungen, eine Checkpoint-Engine und eine oder mehrere Speicherunterstützungsanwendungen umfasst, die in den ein oder mehreren Speichergeräten gespeichert sind und mit jeweiligen Hypervisoren ausgeführt werden.
17. System nach Anspruch 1, wobei Checkpoints mittels der Hardware-Verbindung übertragen werden.
18. Verfahren zum Steuern eines Betriebssystems in Echtzeit, das Folgendes beinhaltet: Empfangen von Gerätedaten an einem primären Knoten und einem sekundären Knoten von einem oder mehreren Geräten in elektrischer Kommunikation mit einem zeitschlitzbasierten Bus, wobei der Bus die Aufgabe hat, Daten gemäß Zeitschlitzes von zwei oder mehr Prioritätsstufen zu empfangen; Verarbeiten der empfangenen Gerätedaten gleichzeitig an dem primären Knoten und dem sekundären Knoten; Erzeugen, während eines ersten Zeitschlitzes, eines Check-

points an dem primären Knoten und Übertragen der Checkpoint-Daten zum sekundären Knoten, wobei der erste Zeitschlitz ein Zeitschlitz niedrigerer Priorität ist, wobei der Kommunikationszeitschlitz niedrigerer Priorität einen Nicht-Echtzeit-Zeitschlitz umfasst; Erzeugen von Ausgabedaten an dem primären Knoten und dem sekundären Knoten als Reaktion auf die Verarbeitung der empfangenen Gerätedaten; und Übertragen der Ausgabedaten von dem primären Knoten über den Bus unter Verwendung eines zweiten Zeitschlitzes, wobei der zweite Zeitschlitz ein Zeitschlitz höherer Priorität ist, wobei sowohl der erste Knoten als auch der zweite Knoten in einem aktiven Modus sind.

19. Verfahren nach Anspruch 18, das ferner das Gating von Netzwerkdaten mit Hilfe eines jeweiligen Netzwerksoftware-Moduls umfasst, das sowohl auf dem primären als auch auf dem sekundären Knoten ausgeführt wird, so dass Ausgabedaten von dem Knoten freigegeben werden, dessen Zustand sich zum aktuellen Checkpoint fortgesetzt hat.
20. Verfahren nach Anspruch 19, das ferner das Synchronisieren der Übertragung von Ausgabedaten mit Zeitschlitzten hoher Priorität umfasst, wobei die Ausgabedaten von dem Betriebssystem empfangen werden.
21. Verfahren nach Anspruch 18 oder Anspruch 19, das ferner das Übertragen von Daten während eines Nicht-Checkpoint-Teils eines Zeitschlitzes niedrigerer Priorität beinhaltet.
22. Verfahren nach Anspruch 18, das ferner das Übertragen, unter Verwendung des primären Knotens, aller seiner jeweiligen ausgehenden Daten zu dem sekundären Knoten beinhaltet, um zu vermeiden, dass solche ausgehenden Daten beim Auftreten eines Ausfalls des primären Knotens verworfen werden.

Revendications

1. Un système informatique tolérant aux pannes en temps réel (7b) comprenant :

un premier dispositif informatique (8a) comprenant

un ou plusieurs dispositifs à mémoire,
un ou plusieurs processeurs,
une première interface réseau conçue de façon à recevoir des données de dispositif et à transmettre des données en sortie par l'intermédiaire d'un bus à base d'intervalle de temps, où les données en sortie sont gé-

nérees à partir du traitement de données de dispositif, et où le bus est conçu de façon à recevoir des données en fonction d'intervalles de temps de deux ou plus niveaux de priorité, et

un premier moteur de point de contrôle en temps réel,
un deuxième dispositif informatique (8b) comprenant

un ou plusieurs dispositifs à mémoire,
un ou plusieurs processeurs,
une deuxième interface réseau conçue de façon à recevoir des données de dispositif par l'intermédiaire du bus à base d'intervalle de temps, et
un deuxième moteur de point de contrôle en temps réel, et

une interconnexion matérielle conçue de façon à permettre un échange de données entre le premier dispositif informatique et le deuxième dispositif informatique,

où le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) traitent simultanément des données de dispositif reçues, et

où des points de contrôle sont générés par le premier ou le deuxième moteur de point de contrôle en temps réel au cours d'intervalles de temps de communication à plus faible priorité attribués sur le bus à base d'intervalle de temps de façon à éviter d'interférer avec toute communication en temps réel de ou vers les premier et deuxième dispositifs informatiques, où les intervalles de temps de communication à plus faible priorité comprennent des intervalles de temps de communication non en temps réel.

2. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 1 où, si le premier dispositif informatique (8a) rencontre une défaillance, le deuxième dispositif informatique poursuit le traitement de données de dispositif de façon à générer les données en sortie et transmettre les données en sortie au moyen du bus à base d'intervalle de temps.

3. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2 où des opérations de point de contrôle sont toujours reportées jusqu'à l'occurrence d'un intervalle de temps de communication à plus faible priorité, même après la production par le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) de données en sortie divergentes avant l'occurrence de l'intervalle de temps de communication à plus faible

priorité.

4. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 3 où le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) retardent tous les deux la libération de données en sortie vers un ou plusieurs dispositifs en communication avec le bus à base d'intervalle de temps jusqu'après un point de contrôle subséquent de façon à éviter la libération de données en sortie divergentes si une défaillance devait se produire. 5
5. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où le premier dispositif informatique (8a) génère des points de contrôle sur une base par cycle lorsque des données en sortie comprenant des données à faible priorité sont transmises au moyen d'un intervalle de temps à faible priorité. 10
6. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où le dispositif est un capteur et les données en sortie comprennent des données de capteur, les données de capteur étant utilisables en tant qu'entrée destinée à un système de commande. 15
7. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 1, où le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) traitent tous les deux activement des données de dispositif et produisent des données en sortie de sorte que, si le premier dispositif informatique (8a) rencontre un ou plusieurs modes de défaillance, un retard de la poursuite du traitement par un deuxième dispositif de traitement (8b) est nul. 20
8. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où la durée d'un ou de plusieurs intervalles de temps est augmentée d'une période de détection, où la période de détection varie d'environ 1 microseconde à environ 100 microsecondes. 25
9. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) retardent tous les deux la libération de données en sortie vers des dispositifs en communication avec le bus à base d'intervalle de temps jusqu'après le point de contrôle suivant de façon à éviter la libération de données en sortie incohérentes, et éventuellement où, si le premier dispositif informatique (8a) fonctionne normalement après un point de contrôle, le premier dispositif informatique (8a) libère toutes données en sortie re- 30

tardées.

10. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où le premier dispositif informatique (8a) transmet la totalité de ses données en sortie au deuxième dispositif informatique (8b) de façon à éviter que de telles données en sortie soient rejetées en cas d'occurrence d'une défaillance du premier dispositif informatique (8a), et éventuellement où, si le premier dispositif informatique (8a) tombe en panne, le deuxième dispositif informatique (8b) reprend en charge la libération des données en sortie du premier dispositif informatique qui ont été générées au cours du dernier point de contrôle. 35
11. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où uniquement le premier dispositif informatique (8a) gère la libération de données en sortie lorsqu'aucune défaillance ne s'est produite et en parallèle partage de telles données en sortie avec le deuxième dispositif informatique (8b). 40
12. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 1, où les données de dispositif provenant d'un ou de plusieurs dispositifs sont transmises au premier dispositif informatique (8a) et au deuxième dispositif informatique (8b) par l'intermédiaire du bus à base d'intervalle de temps. 45
13. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) reçoivent des données entrantes directement par l'intermédiaire d'une radio-diffusion ou d'une multidiffusion. 50
14. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où chaque dispositif parmi le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) comprend un système d'exploitation (OS) non virtuel respectif comprenant un pilote de réseau, un joint de point de contrôle et un pilote d'espace mémoire respectifs conservés en mémoire dans les un ou plusieurs dispositifs à mémoire et exécutés sur chacun des un ou plusieurs processeurs, et éventuellement où le pilote de réseau comprend un réplicateur en communication avec un coordinateur de libération et un notificateur de libération. 55
15. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 3, où un ou plusieurs programmes de mise en réseau conservés en mémoire dans les un ou plusieurs dispositifs à mémoire des premier et deuxième dispositifs infor-

matiques retardent et coordonnent la transmission de données en sortie au moyen d'un ou de plusieurs éléments parmi l'interconnexion matérielle, la première interface réseau et la deuxième interface réseau.

16. Le système informatique tolérant aux pannes en temps réel (7b) selon la Revendication 2, où chaque dispositif parmi le premier dispositif informatique (8a) et le deuxième dispositif informatique (8b) comprend un système d'exploitation (OS) virtuel respectif comprenant une ou plusieurs applications de soutien au réseau respectives, un moteur de point de contrôle et une ou plusieurs applications de soutien à la conservation en mémoire conservées en mémoire dans les un ou plusieurs dispositifs à mémoire et exécutées avec des hyperviseurs respectifs. 10
17. Le système selon la Revendication 1, où des points de contrôle sont transmis au moyen de l'interconnexion matérielle. 20
18. Un procédé de commande d'un système opérationnel en temps réel comprenant la réception de données de dispositif, au niveau d'un noeud primaire et d'un noeud secondaire, à partir d'un ou de plusieurs dispositifs en communication électrique avec un bus à base d'intervalle de temps, où le bus est conçu de façon à recevoir des données en fonction d'intervalles de temps de deux ou plus niveaux de priorité, le traitement des données de dispositif reçues simultanément au niveau du noeud primaire et du noeud secondaire, au cours d'un premier intervalle de temps, la génération d'un point de contrôle au niveau du noeud primaire et la transmission des données de point de contrôle au noeud secondaire, où le premier intervalle de temps est un intervalle de temps à priorité plus faible, où l'intervalle de temps de communication à priorité plus faible comprend un intervalle de temps non en temps réel, la génération de données en sortie au niveau du noeud primaire et du noeud secondaire en réponse au traitement des données de dispositif reçues, et la transmission, à partir du noeud primaire, des données en sortie par l'intermédiaire du bus au moyen d'un deuxième intervalle de temps, où le deuxième intervalle de temps est un intervalle de temps à priorité plus élevée, où à la fois le premier noeud et le deuxième noeud sont dans un mode actif. 25 30 35 40 45 50
19. Le procédé selon la Revendication 18 comprenant en outre le déclenchement périodique de données de réseau au moyen d'un module logiciel de réseau respectif exécuté sur à la fois le noeud primaire et le noeud secondaire de sorte que des données en sortie soient libérées du noeud dont l'état a été transporté vers l'avant vers le point de contrôle actuel. 55

20. Le procédé selon la Revendication 19 comprenant en outre la synchronisation de la transmission de données en sortie avec des intervalles de temps de haute priorité, où les données en sortie sont reçues par le système d'exploitation. 5

21. Le procédé selon la Revendication 18 ou 19 comprenant en outre la transmission de données au cours d'une partie d'un non-point de contrôle d'un intervalle de temps à priorité plus faible.

22. Le procédé selon la Revendication 18 comprenant en outre la transmission, au moyen du noeud primaire, de la totalité de ses données sortantes respectives au noeud secondaire de façon à éviter à telles données sortantes d'être rejetées en cas d'occurrence d'une défaillance du noeud primaire.




FIG. 1A




FIG. 1B

20




FIG. 2

30




FIG. 3




FIG. 4




FIG. 5




FIG. 6





FIG. 7






FIG 9




FIG. 10




FIG. 11




FIG. 12A




FIG. 12B

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 2007174484 A [0002]
- WO 2016054770 A [0002]
- WO 2012084839 A [0002]
- WO 16900909 A [0025]